arXiv:1412.6749v1 [cs.CV] 21 Dec 2014

SENNS: Sparse Extraction Neural NetworkS for Feature Extraction.

Abdulrahman Oladipupo Ibraheem
rahmanoladi@yahoo.com

Computing and Intelligent Systems Research Group
Department of Computer Science and Engineering
Obafemi Awolowo University, Ile-Ife, Nigeria.

December 23, 2014

Abstract

The feature extraction problem occupies a central position in pattern recognition and machine learning. In this concept paper, drawing on ideas from optimisation theory, artificial neural networks (ANN), graph embeddings and sparse representations, I develop a novel technique, termed SENNS (Sparse Extraction Neural NetworkS), aimed at addressing the feature extraction problem. The proposed method uses (preferably deep) ANNs for projecting input attribute vectors to an output space wherein pairwise distances are maximized for vectors belonging to different classes, but minimized for those belonging to the same class, while simultaneously enforcing sparsity on the ANN outputs. The vectors that result from the projection can then be used as features in any classifier of choice. Mathematically, I formulate the proposed method as the minimisation of an objective function which can be interpreted, in the ANN output space, as a *negative factor* of the sum of the squares of the pair-wise distances between output vectors belonging to different classes, added to a *positive factor* of the sum of squares of the pair-wise distances between output vectors belonging to the same classes, plus sparsity and weight decay terms. To derive an algorithm for minimizing the objective function via gradient descent, I use the multi-variate version of the chain rule to obtain the partial derivatives of the function with respect to ANN weights and biases, and find that each of the required partial derivatives can be expressed as a sum of six terms. As it turns out, four of those six terms can be computed using the standard back propagation algorithm; the fifth can be computed via a slight modification of the standard backpropagation algorithm; while the sixth one can be computed via simple arithmetic. Finally, I propose experiments on the ARABASE Arabic corpora of digits and letters, the CMU PIE database of faces, the MNIST digits database, and other standard machine learning databases.

1 Introduction

Most pattern recognition systems comprise three key stages: pre-processing, feature extraction and classification stages. Of these three stages, researchers believe that the feature extraction stage is the most critical. For example, in their review paper, the authors of [1] unequivocally wrote: 'Selection of

a feature extraction method is probably the single most important factor in achieving high recognition performance...' Indeed, we agree with the authors of [1], since our own view is that *feature extraction is a commitment, once made might be irreversible by any classifier, however sophisticated*. Hence, it becomes highly paramount to carefully and rigorously study how these 'commitments' should be made — how should features be extracted for optimal accuracies at the classification stage? While researchers have proposed a plethora of methods (e.g. [2], [3], [4], [5], [6], [7], [8], [10], [9], [12], [13]) aimed at answering this question, it appears that a philosophy that should be followed by any good feature extraction method is the one articulated by Dejiver and Kittler in [14]. They said that feature extraction is the problem of 'extracting from the raw data the information which is most relevant for classification purposes, in the sense of minimizing the within-class pattern variability, while enhancing the between-class pattern variability.' Two of the more popular feature extraction techniques that follow this philosophy are the Linear Discriminant Analysis (LDA) [3] and the Marginal Fisher Analysis (MFA) [12], both of which can be viewed as specific examples of a unifiying concept called graph embediings [12]. Inspired by the above philosophy of [14], we herein propose a technique, termed SENNS (Sparse Extraction Neural NetworkS) and pronounced 'SENSE,' for addreesing the feature extraction problem. Like the LDA and MFA, our SENNS can also be viewed, at least partially, via the lens of the graph embeddings concept. Unlike the MFA and LDA, however, SENNS incorporates a mechanism for seeking sparse features, and employs the apparatus of (preferably deep) non-linear artificial neural networks, rather than linear (or kernel, or tensor) projections, for effecting the transformations that result in the sought features. Mathematically, we formulate our method as the minimisation of an objective function. Via rigorous mathematical analysis, we then derive a gradient descent algorithm for minimizing our objective function. Fortunately, it turns out that our algorithm can be expressed in terms of the standard backpropagation procedure, except, as we shall see, for a little tweaking to accomodate L_1 norms. Finally, we plan to test SENNS on standard machine learning datasets such as the ARABASE Arabic corpora of digits and letters [15], the CMU PIE database of faces [16], the MNIST digits database [17], and other standard machine learning databases.

2 Notation

In describing neural networks, I almost entirely follow the notation of Prof. Andrew Ng as in [18]. In Ng's notation, for supervised training mode, the neural network learns from a training data denoted $(x^{(i)}, y^{(i)})$, $i = 1, 2, ..., m$. The neural network proper consists of n_l layers, $L_1, L_2, ..., L_l, ..., L_{n_l}$, and the number of neurons in the l-th layer is denoted s_l. A weight, denoted $W_{ij}^{(l)}$, connects the j-th neuron of layer l with the i-th neuron of layer $l+1$, while a bias, denoted $b_i^{(l)}$, emanates from layer l and enters the i-th neuron of layer $l+1$. The overall function of the i-th neuron in layer l is to compute an activation denoted $a_i^{(l)}$, which is the result of passing the quantity $z_i^{(l)} = \sum_{j=1}^{s_l} w_{ij}^{(l)} a_j^{(l)} + b_i^{(l)}$ through a transfer function such as the sigmoid or tanh function. As a matter of notational expedience, for each j, the definition $a_i^{(1)} = x^{(j)}$ is employed, which means that the ANN input is viewed as an 'activation'for the first layer. Furthermore, $a^{(l)}$ is used to denote the column vector given by $a^{(l)} = (a_1^{(l)}, a_2^{(l)}, ..., a_{s_l}^{(l)})^T$; and a similar notation applies to $b^{(l)}$ and $z^{(l)}$ as well. Similarly, $W^{(l)}$ denotes the column vector obtained via an ordered concatenation of all the weights linking layer l with layer $l+1$ in the network. Finally, W represents a concatenation of all the weights in the ANN, while b represents a concatenation of all the biases.

Herein, I will follow the above notation, except for the following modifications. Firstly, to avoid

confusion, I will use the index, t, placed within square brackets, to label training data: $(x^{[t]}, y^{[t]})$ $t = 1, 2, \ldots, m$. Based on this, I will write $a_i^{(l)[t]}$ to denote the activation from the i-th neuron of the l-th layer when $x^{[t]}$ is applied as input to the ANN, and I will use $a^{(l)[t]}$ to the denote the column vector of activations associated with the l-th layer and t-th input vector: $a_i^{(l)[t]} = (a_1^{(l)[t]}, a_2^{(l)[t]}, ..., a_{s_l}^{(l)[t]})^T$. However, when the layer in question is clear from context, I will simply write $a_i^{[t]}$ instead of $a_i^{(l)[t]}$, and $a^{[t]}$ instead of $a^{(l)[t]}$, in order to achieve a less clumsy notation. Finally, in dealing with partial derivatives, I shall herein write $\frac{\partial f}{\partial a^{[t]}}\Big|_{\hat{a}^{[t]}}$ as a shorthand for $\frac{\partial f}{\partial a^{[t]}}\Big|_{a^{[t]}=\hat{a}^{[t]}}$. Likewise, we shall write $\frac{\partial f}{\partial a^{[t]}}\Big|_{(\hat{a}^{[t]},\hat{a}^{[u]})}$ instead of $\frac{\partial f}{\partial a^{[t]}}\Big|_{(a^{[t]},a^{[u]})=(\hat{a}^{[t]},\hat{a}^{[u]})}$

3 Formulation of Objective Function

Towards formulating the required objective function, I will begin by introducing two 'functions', $\mathbb{C}$ and $\mathbb{D}$ defined over the Cartesian product of input vectors according to:

$$\mathbb{C}(x^{[t]}, x^{[u]}) = \begin{cases} 1 & \text{if } x^{[t]} \text{ and } x^{[u]} \text{ belong to the } \textbf{same} \text{ class} \\ 0 & \text{otherwise} \end{cases} \tag{1}$$

and

$$\mathbb{D}(x^{[t]}, x^{[u]}) = \begin{cases} 1 & \text{if } x^{[t]} \text{ and } x^{[u]} \text{ belong to } \textbf{different} \text{ classes} \\ 0 & \text{otherwise} \end{cases} \tag{2}$$

Further, if $\mathcal{X} = \{x^{[1]}, x^{[2]}, ..., x^{[m]}\}$ denotes the set of all input vectors, and $\mathcal{X}_2$ is the set of all pairs of the form $(x^{[t]}, x^{[u]})$, $t \in \{1, 2, ..., m\}$ and $u \in \{1, 2, ..., m\}$, which can be drawn from $\mathcal{X}$, then we shall use $M_{\mathbb{C}}(\mathcal{X})$ to denote the number of times that function $\mathbb{C}$ outputs 1 when all the pairs in $\mathcal{X}_2$ are passed through it. A similar definition applies to $M_{\mathbb{D}}(\mathcal{X})$ as well. When the training set is clear from context, we will simply write $M_{\mathbb{C}}$ and $M_{\mathbb{D}}$ instead of $M_{\mathbb{C}}(\mathcal{X})$ and $M_{\mathbb{D}}(\mathcal{X})$ respectively. Next, for some non-negative regularisation constants, $\lambda_1, \lambda_2 \in [0, 1]$, we define a function, $\mathbb{S}(x^{[t]}, x^{[u]})$, as follows:

$$\mathbb{S}(x^{[t]}, x^{[u]}) = \begin{cases} \dfrac{\lambda_1}{M_{\mathbb{C}}} & \text{if } x^{[t]} \text{ and } x^{[u]} \text{ belong to the } \textbf{same} \text{ class} \\ -\dfrac{\lambda_2}{M_{\mathbb{D}}} & \text{otherwise} \end{cases} \tag{3}$$

The connection between function $\mathbb{S}$ and the graph embedding framework [12] should at once be clear. Specifically, it should be clear that $\mathbb{S}$ plays the role of the weights on the edges of the graphs underlying graph embeddings. We see that, similar to the Marginal Fisher Analysis (MFA) described in [12], $\mathbb{S}$ connects data points belonging to the same class with a positive weight, but connects those belonging to different classes with a negative weight. However, as we shall soon see, the proposed mathod herein differs from that in [12] in three key ways. First, herein, we effect our projections via ANNs, unlike [12] who employed either linearisation, kernelisation or tensorisation. Second, herein, we impose a sparsity requirement on the sought features, thereby seeking to take advantage of the well known benefits of sparse features; see [19] and [20] for instance. Indeed, we would like to see the effect of the sparsity term on the ability of our gradient descent algorithm to locate a global minimum for our non-convex

objective function. Thirdly, the overall structure of our objective function herein is different from that in [12], since ours is a regularised sum of terms, whereas theirs is a quotient of terms. The above three things also distinguish the method proposed herein from Linear Discriminant Analysis (LDA) [3]. We can now spell out the objective function we wish to minimize as follows:

$$J(W,b) \;=\; \frac{1}{2}\sum_{t=1}^{m}\sum_{u=1}^{m}\mathbb{S}(x^{[t]},x^{[u]})||a^{[t]} \;-\; a^{[u]}||^2 \quad + \quad \frac{\lambda_3}{m}\sum_{t=1}^{m}||a^{[t]}||_1 \;+\; \frac{\lambda_4}{2}\sum_{l=1}^{n_l}\sum_{i=1}^{s_{l+1}}\sum_{j=1}^{s_l}W_{ij}^{(l)2} \quad (4)$$

To achieve a slightly less clumsy notation in the above equation, we have written the activations associated with the ANN's output layer as $a^{[t]}$ instead of $a^{(n_l)[t]}$. We shall carry on this practice henceforth, unless otherwise stated. Further, $||.||_1$ denotes the L_1 norm. Also, notice that the above objective function implicitly incorporates the regularizers, λ_1, λ_2, via the inclusion of $\mathbb{S}(x^{[t]}, x^{[u]})$. In addition, λ_3 and λ_4 are also regularizers such that $\lambda_3, \lambda_4 \in [0,1]$ and $\lambda_1 + \lambda_2 + \lambda_3 + \lambda_4 = 1$. By simply substituting $\mathbb{S}(x^{[t]}, x^{[u]})$ from Equation 3 into Equation 4, we have:

$$J(W,b) = \frac{\lambda_1}{2M_{\mathbb{C}}}\sum_{t=1}^{m}\sum_{u=1}^{m}\mathbb{C}(x^{[t]},x^{[u]})||a^{[t]} \;-\; a^{[u]}||^2 \;-\; \frac{\lambda_2}{2M_{\mathbb{D}}}\sum_{t=1}^{m}\sum_{u=1}^{m}\mathbb{D}(x^{[t]},x^{[u]})||a^{[t]} \;-\; a^{[u]}||^2 + \frac{\lambda_3}{m}\sum_{t=1}^{m}||a^{[t]}||_1 \;+\; \frac{\lambda_4}{2}\sum_{l=1}^{n_l}\sum_{i=1}^{s_{l+1}}\sum_{j=1}^{s_l}W_{ij}^{(l)2} \quad (5)$$

The form in Equation 5 above highlights the 'graph embedding' aspect of our formulation. The first term gives a measure, in the output space of the ANN, of how widely separated output vectors belonging to the same class are. Clearly, we wish to minimize this non-negative quantity. On the contrary, the second term, *excluding its negative sign*, gives a measure, in the output space of the ANN, of how widely separated output vectors belonging to different classes are. We wish to maximize this **non**-negative measure, by minimizing the negative quantity that results when the negative sign is pre-fixed to it. The third term is a sparsity term by which we wish to make the extracted features sparse. The fourth term is a weight decay term which prevents the weights from becoming too large, and helps prevent overfitting. Finally, we see that the parameters, $\lambda_1, \lambda_2, \lambda_3, \lambda_4$, allow us to control the relative amount of significance that the objective function attaches to the four objectives it is trying to achieve.

A question naturally arises pertaining to the computational feasibility of the sums appearing in Equation 5. In particular, as we shall see, the sums involving index variables t and m will carry over directly to the algorithm for minimizing the objective function in Equation 5. This means that the term involving $\mathbb{C}(x^{[t]}, x^{[u]})$ (as well as the term involving $\mathbb{D}(x^{[t]}, x^{[u]})$) in the objective function would require $O(m^2)$ time, which becomes undesirable as m grows. To ameliorate this, we are going to propose two heuristics for the two sums, as follows. We first consider the case of the term involving $\mathbb{D}(x^{[t]}, x^{[u]})$. Upfront, we point out that the heuristic leads to a maximisation, in the ANN output space, of the sum of the distances between each input vector from a given class, and its nearest neighbour from each of the other classes, thereby given rise to a *maximisation of the minimum distance formulation*, reminiscent of support vector machines [**?**]. To proceed, let $\mathcal{T} = \{x^{[1]}, x^{[2]}, \ldots, x^{[m]}\}$, and let there be N classes, denoted $\Omega_1, \Omega_2, \ldots, \Omega_N$ in the classification problem for which we are extracting features. Now, for each $x^{[t]} \in \Omega_p$, we define a set, $\mathbb{D}_{x^{[t]}}$ containing $N-1$ elements as follows: $\mathbb{D}_{x^{[t]}} = \{d_1, d_2, \ldots d_{p-1}, d_{p+1}, \ldots, d_N\}$, such that each d_p belongs to class Ω_p and each d_p is the *nearest* member of Ω_p from $x^{[t]}$. Our heuristic is to replace the quantity $\sum_{t=1}^{m}\sum_{u=1}^{m}\mathbb{D}(x^{[t]}, x^{[u]})||a^{[t]} \;-\; a^{[u]}||^2$ by $\sum\limits_{x^{[t]}\in\mathcal{T}}\sum\limits_{d_q\in\mathbb{D}_{x^{[t]}}}||a^{[t]} \;-\; a^{[q]}||^2$ where $a^{[q]}$

denotes the ANN output-layer vector of activations associated with input vector $d_q \in \mathbb{D}_{x^{[t]}}$. It should be clear that this new sum is $O(\gamma m)$, where $\gamma = N - 1$, compared to the original sum which is $O(m^2)$. Since most problems usually have $m >> N - 1$, we expect this heuristic to lead to significant gains in computational feasibility in most cases.

However, for the case of the term involving $\mathbb{C}(x^{[t]}, x^{[u]})$, we propose an heuristic that leads to a situation wherein the sum of the distances between each input vector and its k-farthest neighbours, all in the same class as the input vector, is minimised. Formally, for each $x^{[t]} \in \Omega_p$, we simply define the set $\mathbb{C}_{x^{[t]}} = \tilde{N}_k(x^{[t]})$, where as usual, $\tilde{N}_k(x^{[t]})$ is the set of the k *farthest* elements from $x^{[t]}$ in Ω_p. In Equation 5, we then replace the quantity $\sum_{t=1}^{m}\sum_{u=1}^{m}\mathbb{C}(x^{[t]}, x^{[u]})||a^{[t]} - a^{[u]}||^2$ by $\sum\limits_{x^{[t]}\in\mathcal{T}}\sum\limits_{d_q\in\mathbb{C}_{x^{[t]}}}||a^{[t]} - a^{[q]}||^2$. Again, in this case, we see that, for most cases, the heuristic can lead to improved computational feasibility since the new sum is $O(km)$, and k can be chosen to be far lesser than m.

Since the above heuristics constitute just an example of a host of possible heuristics that can be applied to alleviate the computational feasibilty issue in Equation 5, we therefore think that it would be better for us to develop our technique for the general case formulated in Equation 5, especially considering the fact that it should be clear how to adapt the developed technique to any particular heuristic of interest.

So now, let us go back to Equation 4 (from which Equation 5 derives). It is expedient to denote the first term in the equation by $J_1(W, b)$, the second term by $J_2(W, b)$, and the third term by $J_3(W)$. Hence, Equation 4 can be re-written in the form: $J(W, b) = J_1(W, b) + J_2(W, b) + J_3(W)$

We now consider how to minimize $J(W, b)$ via gradient descent. A key step is the computation of $\nabla_{W^{(l)}} J(W, b)$ and $\nabla_{b^{(l)}} J(W, b)$, and this distills to the computation of $\frac{\partial J(W,b)}{\partial W_{ij}^{(l)}}$ and $\frac{\partial J(W,b)}{\partial b_i^{(l)}}$, for all $l = 1, 2, ..., n_l - 1$, for all $j = 1, 2, ...s_l$, and for all $i = 1, 2, ..., s_{l+1}$. I shall illustrate my overall approach by showing how to compute $\frac{\partial J(W,b)}{\partial W_{ij}^{(l)}}$, since the computation of $\frac{\partial J(W,b)}{\partial b_i^{(l)}}$ is analogous. To this end, I begin with a rather trivial step and write:

$$\frac{\partial J(W,b)}{\partial W_{ij}^{(l)}} = \frac{\partial J_1(W,b)}{\partial W_{ij}^{(l)}} + \frac{\partial J_2(W,b)}{\partial W_{ij}^{(l)}} + \frac{\partial J_3(W)}{\partial W_{ij}^{(l)}} \tag{6}$$

But, the third partial derivative, $\frac{\partial J_3(W)}{\partial W_{ij}^{(l)}}$, is particularly straightforward to compute: $\frac{\partial J_3(W)}{\partial W_{ij}^{(l)}} = W_{ij}^{(l)}$. Plugging this into Equation 6, we readily obtain:

$$\frac{\partial J(W,b)}{\partial W_{ij}^{(l)}} = \frac{\partial J_1(W,b)}{\partial W_{ij}^{(l)}} + \frac{\partial J_2(W,b)}{\partial W_{ij}^{(l)}} + W_{ij}^{(l)} \tag{7}$$

To proceed, let us introduce two definitions: $\hat{J}_1(W, b) = \frac{1}{2}||a^{[t]} - a^{[u]}||^2$, and $\hat{J}_2(W, b) = ||a^{[t]}||_1$. By comparing Equation 4 with the expression $J(W, b) = J_1(W, b) + J_2(W, b)J_3(W)$, observe that the first definition above allows us to write: $J_1(W, b) = \sum\limits_{t=1}^{m}\sum\limits_{u=1}^{m}\mathbb{S}(x^{[t]}, x^{[u]})\hat{J}_1(W, b)$, so that:

$$\frac{\partial J_1(W,b)}{\partial W_{ij}^{(l)}} = \sum_{t=1}^{m}\sum_{u=1}^{m}\mathbb{S}(x^{[t]}, x^{[u]})\frac{\partial \hat{J}_1(W,b)}{\partial W_{ij}^{(l)}} \tag{8}$$

In the same vein, notice that the second definition above permits us to write: $J_2(W,b) = \frac{\lambda_3}{m}\sum_{t=1}^{m}\hat{J}_2(W,b)$, so that:

$$\frac{\partial J_2(W,b)}{\partial W_{ij}^{(l)}} = \frac{\lambda_3}{m}\sum_{t=1}^{m}\frac{\partial \hat{J}_2(W,b)}{\partial W_{ij}^{(l)}} \tag{9}$$

. In rounding off this section, we should point out that the expression for $\frac{\partial J(W,b)}{\partial W_{ij}^{(l)}}$ given in Equation 7 will be particularly useful in the gradient descent algorithm we shall be proposing in Section 5 for minimizing $J(W,b)$. The algorithm will first use Equations 8 and 9 to compute $\frac{\partial J_1(W,b)}{\partial W_{ij}^{(l)}}$ and $\frac{\partial J_2(W,b)}{\partial W_{ij}^{(l)}}$ respectively; it will then use Equation 7 to compute $\frac{\partial J(W,b)}{\partial W_{ij}^{(l)}}$, before proceeding to update ANN weights in a gradient-descent style. However, it is clear from Equations 8 and 9 that if the algorithm is to be effective, then we must find ways of computing $\frac{\partial \hat{J}_1(W,b)}{\partial W_{ij}^{(l)}}$ and $\frac{\partial \hat{J}_2(W,b)}{\partial W_{ij}^{(l)}}$. The next two sections are devoted to these two tasks.

4 Deriving the Partial Derivatives of $\hat{J}_1(W,b)$, and an Algorithm for Computing them.

In this section, we focus on the computation of $\frac{\partial \hat{J}_1(W,b)}{\partial W_{ij}^{(l)}}$. From the preceding section, we know:

$$\hat{J}_1(W,b) = \frac{1}{2}||a^{[t]} - a^{[u]}||^2 \tag{10}$$

Next, I define vector A to be the column vector formed by stacking the column vector, $a^{[t]}$ atop the column vector $a^{[u]}$. That is, $A = ({a^{[t]}}^T, {a^{[u]}}^T)^T$, where ${a^{[t]}}^T$ denotes the transpose of $a^{[t]}$. Notice that $\hat{J}_1(W,b)$ is a function of A, so that one may speak of computing $\nabla_A \hat{J}_1(W,b)$, $\nabla_{a^{[t]}}\hat{J}_1(W,b)$ and $\nabla_{a^{[u]}}\hat{J}_1(W,b)$. Also, observe that, since each of $a^{[t]}$ and $a^{[u]}$ is a function of ANN weights, W, and biases, b, it follows that A is also a function of weights and biases, so that one can as well speak of computing $\frac{\partial A}{\partial W_{ij}^{(l)}}$. Now, going back to Equation 10, we can now use the chain rule to write an expression for $\frac{\partial \hat{J}_1(W,b)}{\partial W_{ij}^{(l)}}$:

$$\frac{\partial \hat{J}_1(W,b)}{\partial W_{ij}^{(l)}} = \left(\frac{\partial A}{\partial W_{ij}^{(l)}}\right)^T \nabla_A \hat{J}_1(W,b) \tag{11}$$

Being a dot product of two vectors, the expression above is a sum of terms. We can break the sum into two parts, one associated with the column vector, $a^{[t]}$ and the other associated with the column vector $a^{[u]}$:

$$\frac{\partial \hat{J}_1(W,b)}{\partial W_{ij}^{(l)}} = \left(\frac{\partial a^{[t]}}{\partial W_{ij}^{(l)}}\right)^T \nabla_{a^{[t]}}\hat{J}_1(W,b) + \left(\frac{\partial a^{[u]}}{\partial W_{ij}^{(l)}}\right)^T \nabla_{a^{[u]}}\hat{J}_1(W,b) \tag{12}$$

I will now go down a path to argue that each of the terms in Equation 12 above can be computed using the standard backpropagation algorithm. To this end, I consider a function that has a very similar form to $\hat{J}_1(W,b)$, except for a very small difference. In particular, I consider the function $\tilde{B}(a^{[t]};\ c^{[t]})$, which I herein term a *Back-Propagatable* function, and which I define according to:

$$\tilde{B}(a^{[t]};\ c^{[t]}) = \frac{1}{2}||a^{[t]}\ - c^{[t]}||^2 \tag{13}$$

where, $c^{[t]}$ is a constant that is independent of the ANN's weights and biases. We ask the reader to compare the function $\tilde{B}(a^{[t]};\ c^{[t]}) = \frac{1}{2}||a^{[t]}\ - c^{[t]}||^2$ with the function $\hat{J}_1(W,b) = \frac{1}{2}||a^{[t]}\ - a^{[u]}||^2$, defined in Equation 10. In particular, the reader should note that $\tilde{B}(a^{[t]};\ c^{[t]})$ can be obtained from $E(W,b;\ x^{[t]},x^{[u]})$ simply by replacing the *function* (of weights and biases) $a^{[u]}$ in the latter by the *constant* $c^{[t]}$. However, what is more important is that the back-propagatable function in Equation 13 above plays a central role in the expression of the sum of squares error that an ANN aimed at classification must try to minimize in supervised learning mode. Given training data, $\mathcal{T} = (x^{[t]}, y^{[t]}), t = 1, 2, 3, ..., m$, I can recall that that sum of squares error can be written as:

$$B(W,b) = \frac{1}{2m}\sum_{t=1}^{m}||a^{[t]}\ - y^{[t]}||^2 \tag{14}$$

Clearly, in the above equation, $y^{[t]}$ is a constant independent of ANN weights and biases. Moreover, it is also clear that the expression inside the summation on the right hand side of the equation perfectly fits into:

$$\tilde{B}(a^{[t]};\ y^{[t]}) = \frac{1}{2}||a^{[t]}\ - y^{[t]}||^2 \tag{15}$$

To proceed, we put Equation 15 into Equation 14 and obtain:

$$B(W,b) = \frac{1}{m}\sum_{t=1}^{m}\tilde{B}(a^{[t]};\ y^{[t]}) \tag{16}$$

In supervised learning mode, the objective of the ANN is to minimize the total error, $B(W,b)$.One of the most frequently employed techniques for minimizing $B(W,b)$ is the gradient descent algorithm, and this requires the partial derivatives of $B(W,b)$ with respect to weights and biases. Now, the partial derivative of $B(W,b)$ with respect to an arbitrary weight $W_{ij}^{(l)}$ can be expressed as:

$$\frac{\partial B(W,b)}{\partial W_{ij}^{(l)}} = \frac{1}{m}\sum_{t=1}^{m}\frac{\partial \tilde{B}(a^{[t]};\ y^{[t]})}{\partial W_{ij}^{(l)}} \tag{17}$$

So in essence, the task of computing $\frac{\partial B(W,b)}{\partial W_{ij}^{(l)}}$ boils down to the computation of $\frac{\partial \tilde{B}(a^{[t]};\ y^{[t]})}{\partial W_{ij}^{(l)}}$, and then summing over all $t \in \{1, 2, ..., m\}$ (i.e. summing over the training data); *the backpropagation algorithm is normally employed for the computation of* $\frac{\partial \tilde{B}(a^{[t]};\ y^{[t]})}{\partial W_{ij}^{(l)}}$. *Generalizing, we see that the backpropagation algorithm can always be used to compute the partial derivatives,* $\frac{\partial \tilde{B}(a^{[t]};\ c^{[t]})}{\partial W_{ij}^{(l)}}$, ***of any***

function of the form $\tilde{B}(a^{[t]};\ c^{[t]}) = \frac{1}{2}||a^{[t]} - c^{[t]}||^2$, *such that* $c^{[t]}$ *is a constant* which plays the role which the target output value, $y^{[t]}$, plays in the supervised learning mode of ANNs. To proceed from here, we shall let $\hat{a}^{[u]}$ denote the specific constant value that the function $a^{[u]}$ ($a^{[u]}$ is a function of ANN weights and biases) evaluates to for a given value of u and a given set of weights and biases. We then consider the function $\tilde{B}(a^{[t]};\ \hat{a}^{[u]})$, in which $\hat{a}^{[u]}$ is playing the role that the target output value, $y^{[t]}$, plays in the function $\tilde{B}(a^{[t]};\ y^{[t]})$. With the foregoing in mind, let us now try to write an expression for $\frac{\partial \tilde{B}(a^{[t]};\ \hat{a}^{[u]})}{\partial W_{ij}^{(l)}}$. With the aid of the chain rule, we have:

$$\frac{\partial \tilde{B}(a^{[t]};\ \hat{a}^{[u]})}{\partial W_{ij}^{(l)}} \ = \ \Big(\frac{\partial a^{[t]}}{\partial W_{ij}^{(l)}}\Big)^T \nabla_{a^{[t]}} \tilde{B}(a^{[t]};\ \hat{a}^{[u]}) \tag{18}$$

We now try to compare the quantity, $\Big(\frac{\partial a^{[t]}}{\partial W_{ij}^{(l)}}\Big)^T \nabla_{a^{[t]}} \tilde{B}(a^{[t]};\ \hat{a}^{[u]})$ on the right hand side of Equation 18 above with the quantity, $\Big(\frac{\partial a^{[t]}}{\partial W_{ij}^{(l)}}\Big)^T \nabla_{a^{[t]}} \hat{J}_1(W,b)$, which occurs as the first term on the right hand side of Equation 12. We claim that, for a given value, $\hat{a}^{[u]}$, both are equal, in the sense that: $\Big[\Big(\frac{\partial a^{[t]}}{\partial W_{ij}^{(l)}}\Big)^T \nabla_{a^{[t]}} \hat{J}_1(W,b)\Big]\Big|_{\hat{a}^{[u]}} = \Big(\frac{\partial a^{[t]}}{\partial W_{ij}^{(l)}}\Big)^T \nabla_{a^{[t]}} \tilde{B}(a^{[t]};\ \hat{a}^{[u]})$, where, as explained in the section on notation, we have written $\Big[\Big(\frac{\partial a^{[t]}}{\partial W_{ij}^{(l)}}\Big)^T \nabla_{a^{[t]}} \hat{J}_1(W,b)\Big]\Big|_{\hat{a}^{[u]}}$ instead of $\Big[\Big(\frac{\partial a^{[t]}}{\partial W_{ij}^{(l)}}\Big)^T \nabla_{a^{[t]}} \hat{J}_1(W,b)\Big]\Big|_{a^{[u]}=\hat{a}^{[u]}}$. Before we show the equality, let us first highlight what we stand to gain if they are truly equal. In particular, $\Big[\Big(\frac{\partial a^{[t]}}{\partial W_{ij}^{(l)}}\Big)^T \nabla_{a^{[t]}} \hat{J}_1(W,b)\Big]\Big|_{\hat{a}^{[u]}} = \Big(\frac{\partial a^{[t]}}{\partial W_{ij}^{(l)}}\Big)^T \nabla_{a^{[t]}} \tilde{B}(a^{[t]};\ \hat{a}^{[u]})$ clearly implies $\Big[\Big(\frac{\partial a^{[t]}}{\partial W_{ij}^{(l)}}\Big)^T \nabla_{a^{[t]}} \hat{J}_1(W,b)\Big]\Big|_{\hat{a}^{[u]}} = \frac{\partial \tilde{B}(a^{[t]};\ \hat{a}^{[u]})}{\partial W_{ij}^{(l)}}$. But, according to our discussion in the previous paragraph, we know that $\frac{\partial \tilde{B}(a^{[t]};\ \hat{a}^{[u]})}{\partial W_{ij}^{(l)}}$ can be computed via backpropagation. Thus, if the equality is really true, we would have in essence found a means of computing $\Big[\Big(\frac{\partial a^{[t]}}{\partial W_{ij}^{(l)}}\Big)^T \nabla_{a^{[t]}} \hat{J}_1(W,b)\Big]\Big|_{\hat{a}^{[u]}}$. Now, this would be very important to us because, as given in Equation 12, $\Big(\frac{\partial a^{[t]}}{\partial W_{ij}^{(l)}}\Big)^T \nabla_{a^{[t]}} \hat{J}_1(W,b)$ is one of the two terms involved in our expression for $\frac{\partial \hat{J}_1(W,b)}{\partial W_{ij}^{(l)}}$, which in turn is a partial derivative which we need for minimizing our objective function via gradient descent. Indeed, it is *very simple* to show the required equality. Specifically, to show $\Big[\Big(\frac{\partial a^{[t]}}{\partial W_{ij}^{(l)}}\Big)^T \nabla_{a^{[t]}} \hat{J}_1(W,b)\Big]\Big|_{\hat{a}^{[u]}} = \Big(\frac{\partial a^{[t]}}{\partial W_{ij}^{(l)}}\Big)^T \nabla_{a^{[t]}} \tilde{B}(a^{[t]};\ \hat{a}^{[u]})$, all we need show is $\nabla_{a^{[t]}} \hat{J}_1(W,b)\big|_{\hat{a}^{[u]}} =$

$\nabla_{a^{[t]}}\tilde{B}(a^{[t]};\ \hat{a}^{[u]})$, since $\Big(\dfrac{\partial a^{[t]}}{\partial W_{ij}^{(l)}}\Big)^T\Big|_{\hat{a}^{[u]}} = \Big(\dfrac{\partial a^{[t]}}{\partial W_{ij}^{(l)}}\Big)^T$. A simple 'proof' follows thus:

'Proof' *of* $\nabla_{a^{[t]}}\hat{J}_1(W,b)\big|_{\hat{a}^{[u]}} = \nabla_{a^{[t]}}\tilde{B}(a^{[t]};\ \hat{a}^{[u]})$.

We begin with $\hat{J}_1(W,b) = \dfrac{1}{2}||a^{[t]} - a^{[u]}||^2 = \dfrac{1}{2}\sum_{i=1}^{s_{n_l}}(a_i^{[t]} - a_i^{[u]})^2$, where $a_i^{[t]}$ denotes the activation of the i-th neuron in the output layer when the training example $x^{[t]}$ is applied as input to the ANN, and s_{n_l} is the number of neurons in the output layer. We then find $\nabla_{a^{[t]}}\hat{J}_1(W,b) = (a_1^{[t]} - a_1^{[u]},\ a_2^{[t]} - a_2^{[u]},\ \ldots,\ a_{s_{n_l}}^{[t]} - a_{s_{n_l}}^{[u]})$. Hence, for a given output value $\hat{a}^{[u]}$ of function $a^{[u]}$ we obtain $\nabla_{a^{[t]}}\hat{J}_1(W,b)\big|_{\hat{a}^{[u]}} = (a_1^{[t]} - \hat{a}_1^{[u]},\ a_2^{[t]} - \hat{a}_2^{[u]},\ \ldots,\ a_{s_{n_l}}^{[t]} - \hat{a}_{s_{n_l}}^{[u]})$. On the other hand, we have $\tilde{B}(a^{[t]};\ \hat{a}^{[u]}) = ||a^{[t]} - \hat{a}^{[u]}||^2 = \sum_{i=1}^{s_{n_l}}\dfrac{1}{2}(a_i^{[t]} - \hat{a}_i^{[u]})^2$, from which we may compute: $\nabla_{a^{[t]}}\tilde{B}(a^{[t]};\ \hat{a}^{[u]}) = (a_1^{[t]} - \hat{a}_1^{[u]},\ a_2^{[t]} - \hat{a}_2^{[u]},\ \ldots\ ,\ a_{s_{n_l}}^{[t]} - \hat{a}_{s_{n_l}}^{[u]})$. Hence, $\nabla_{a^{[t]}}\hat{J}_1(W,b)\big|_{\hat{a}^{[u]}} = \nabla_{a^{[t]}}\tilde{B}(a^{[t]};\ \hat{a}^{[u]})$, as required. □

Before moving on, let us quickly 'save'a key implication of the above result for future reference:

$$\Big[\Big(\frac{\partial a^{[t]}}{\partial W_{ij}^{(l)}}\Big)^T\nabla_{a^{[t]}}\hat{J}_1(W,b)\Big]\Big|_{\hat{a}^{[u]}} = \frac{\partial\tilde{B}(a^{[t]};\ \hat{a}^{[u]})}{\partial W_{ij}^{(l)}} \tag{19}$$

At this juncture, we pause to take stock of what we have achieved so far. Essentially, what we have achieved up to this point is a means for computing the first of the two terms required for the computation of $\dfrac{\partial\hat{J}_1(W,b)}{\partial W_{ij}^{(l)}}$ as given in Equation 12. Hence, my natural next line of action is to consider how to compute the second of the two terms, which is the quantity $\Big(\dfrac{\partial a^{[u]}}{\partial W_{ij}^{(l)}}\Big)^T\nabla_{a^{[u]}}\hat{J}_1(W,b)$. As it turns out, the approach to the computation of $\Big(\dfrac{\partial a^{[u]}}{\partial W_{ij}^{(l)}}\Big)^T\nabla_{a^{[u]}}\hat{J}_1(W,b)\Big|_{\hat{a}^{[t]}}$ is perfectly analogous to that of $\Big(\dfrac{\partial a^{[t]}}{\partial W_{ij}^{(l)}}\Big)^T\nabla_{a^{[t]}}\hat{J}_1(W,b)\big|_{\hat{a}^{[u]}}$, which have already dealt with, so I will not go into all the details. Only two things are required. First, one needs to take advantage of the fact that $(-1)^2 = 1$, which allows one to re-write $\hat{J}_1(W,b)$ in the form $\hat{J}_1(W,b) = \dfrac{1}{2}||a^{[u]} - a^{[t]}||^2$. Second, one considers a back-propagatable function of the form $\tilde{B}(a^{[u]};\ \hat{a}^{[t]}) = \dfrac{1}{2}||a^{[u]} - \hat{a}^{[t]}||^2$ (Compare this with the form $\tilde{B}(a^{[t]};\ \hat{a}^{[u]}) = \dfrac{1}{2}||a^{[t]} - \hat{a}^{[u]}||^2$ which we had used earlier on). Based on this, the current scenario becomes perfectly analogous to the scenario we had earlier dealt with earlier on, while trying to show the claim, $\Big[\Big(\dfrac{\partial a^{[t]}}{\partial W_{ij}^{(l)}}\Big)^T\nabla_{a^{[t]}}\hat{J}_1(W,b)\Big]\Big|_{\hat{a}^{[u]}} = \Big(\dfrac{\partial a^{[t]}}{\partial W_{ij}^{(l)}}\Big)^T\nabla_{a^{[t]}}\tilde{B}(a^{[t]};\ \hat{a}^{[u]})$. Hence, with the analogy,

it should be clear that, for the present scenario, we must have: $\left[\left(\frac{\partial a^{[u]}}{\partial W_{ij}^{(l)}}\right)^T \nabla_{a^{[u]}} \hat{J}_1(W,b)\right]\Big|_{\hat{a}^{[t]}} = \left(\frac{\partial a^{[u]}}{\partial W_{ij}^{(l)}}\right)^T \nabla_{a^{[u]}} \tilde{B}(a^{[u]};\ \hat{a}^{[t]})$. Furthermore, for the present situation, we have an analog of Equation 19 as follows:

$$\left[\left(\frac{\partial a^{[u]}}{\partial W_{ij}^{(l)}}\right)^T \nabla_{a^{[u]}} \hat{J}_1(W,b)\right]\Big|_{\hat{a}^{[t]}} = \frac{\partial \tilde{B}(a^{[u]};\ \hat{a}^{[t]})}{\partial W_{ij}^{(l)}} \tag{20}$$

Now, when we combine Equations 12, 19, and 20, it should not be hard to see that, for a given pair of constants $(\hat{a}^{[t]}, \hat{a}^{[u]})$, we must have:

$$\frac{\partial \hat{J}_1(W,b)}{\partial W_{ij}^{(l)}}\Big|_{(\hat{a}^{[t]},\hat{a}^{[u]})} = \frac{\partial \tilde{B}(a^{[t]};\ \hat{a}^{[u]})}{\partial W_{ij}^{(l)}}\Big|_{\hat{a}^{[t]}} + \frac{\partial \tilde{B}(a^{[u]};\ \hat{a}^{[t]})}{\partial W_{ij}^{(l)}}\Big|_{\hat{a}^{[u]}} \tag{21}$$

A key essence of Equation 21 above is that it provides an avenue for computing $\frac{\partial \hat{J}_1(W,b)}{\partial W_{ij}^{(l)}}$, for a given pair of constants $(\hat{a}^{[t]}, \hat{a}^{[u]})$, using the backpropagation algorithm. Let us illustrate the 'avenue' by considering how to compute the quantity $\frac{\partial \tilde{B}(a^{[t]};\ \hat{a}^{[u]})}{\partial W_{ij}^{(l)}}\Big|_{\hat{a}^{[t]}}$, which occurs in the Equation. To compute $\frac{\partial \tilde{B}(a^{[t]};\ \hat{a}^{[u]})}{\partial W_{ij}^{(l)}}\Big|_{\hat{a}^{[t]}}$, we simply call the standard backpropagation procedure, passing the current training example, $\hat{a}^{[u]}$ and $x^{[t]}$. In the procedure, $\hat{a}^{[u]}$ plays the role of the target output vector; while $x^{[t]}$ is used to calculate $\hat{a}^{[t]}$, which then plays the role of the current vector of output-layer activations from the ANN. On the flip side, to compute $\frac{\partial \tilde{B}(a^{[u]};\ \hat{a}^{[t]})}{\partial W_{ij}^{(l)}}\Big|_{\hat{a}^{[u]}}$, we pass the training example, $x^{[u]}$ (not $x^{[t]}$), along with the pair, $\hat{a}^{[t]}$ (not $\hat{a}^{[u]}$), as arguments to the standard backpropagation procedure. This time, $\hat{a}^{[u]}$, which can be computed from $x^{[u]}$, plays the role of the current vector of activations produced by the ANN, while $\hat{a}^{[t]}$ plays the role of the target output vector. Based on the foregoing, we naturally derive the following algorithm for computing $\frac{\partial \hat{J}_1(W,b)}{\partial W_{ij}^{(l)}}$:

Algorithm 1 Harnessesing the Standard Back-Propagation (SBP) Procedure for Computing $\dfrac{\partial \hat{J}_1(W,b)}{\partial W_{ij}^{(l)}}$

Inputs: A set of current ANN weights, $W_{ij}^{(l)}$, and biases, $b_i^{(l)}$, $l = 1, 2, \ldots, n_l$, $i = 1, 2, \ldots, s_l$, and $j = 1, 2, \ldots, s_{l-1}$, along with a pair of input vectors $(x^{[t]}, x^{[u]})$.

1). Use the input vector, $x^{[t]}$, to perform an initialisation by filling ANN 'input activations' with $x^{[t]}$.

That is, in the input layer, FOR each $i = 1, 2, \ldots, s_1$, set:
$a_i^{(1)} := x_i^{[t]}$
END FOR

2). Perform a Feedforward pass through the ANN, computing all activations through out the network.

That is, FOR each $l = 1, 2, \ldots, n_l - 1$, FOR each $i = 1, 2, \ldots, s_{l+1}$, FOR each $j = 1, 2, \ldots, s_l$, set:

2a). $z_i^{(l+1)[t]} := W_{ij}^{(l)} a_i^{(l)} + b_i^{(l)}$
2b). $a_i^{(l+1)[t]} := f(z_i^{(l+1)[t]})$
END FOR, END FOR, END FOR

3). Obtain $\hat{a}^{[t]}$ as the vector of output layer activations computed in STEP 2b above.
That is, FOR each $i = 1, 2, \ldots, s_{n_l}$, set:
$\hat{a}_i^{[t]} := a_i^{(n_l)[t]}$.
END FOR

4). Repeat STEPS 1 to 3, but this time using $x^{[u]}$, rather than $x^{[t]}$ (and using variable u rather than t), thereby obtaining vector $\hat{a}^{[u]}$.

5). Use the SBP procedure to compute $\dfrac{\partial \tilde{B}(a^{[t]};\, \hat{a}^{[u]})}{\partial W_{ij}^{(l)}}\Bigg|_{\hat{a}^{[t]}}$, by passing $x^{[t]}$ and $\hat{a}^{[u]}$ to the procedure; this first time, $x^{[t]}$ plays the role of the current ANN input, while $\hat{a}^{[u]}$ plays the role of the target output vector:

$$\Delta\tilde{B}_1 := \frac{\partial \tilde{B}(a^{[t]};\, \hat{a}^{[u]})}{\partial W_{ij}^{(l)}}\Bigg|_{\hat{a}^{[t]}}$$

6). Use the SBP procedure to compute $\dfrac{\partial \tilde{B}(a^{[u]};\, \hat{a}^{[t]})}{\partial W_{ij}^{(l)}}\Bigg|_{\hat{a}^{[u]}}$, by passing $x^{[u]}$ and $\hat{a}^{[t]}$ to the procedure; this second time, $x^{[u]}$ plays the role of the current ANN input, while $\hat{a}^{[t]}$ plays the role of the target output vector:

$$\Delta\tilde{B}_2 := \frac{\partial \tilde{B}(a^{[u]};\, \hat{a}^{[t]})}{\partial W_{ij}^{(l)}}\Bigg|_{\hat{a}^{[u]}}$$

7). Compute the final result:

$$\frac{\partial \hat{J}_1(W,b)}{\partial W_{ij}^{(l)}} := \Delta\tilde{B}_1 + \Delta\tilde{B}_2$$

5 Deriving the Partial Derivatives of $\frac{\partial \hat{J}_2(W,b)}{\partial W_{ij}^{(l)}}$ and an Algorithm for Computing them.

In the preceding section, we derived an expression, and corresponding algorithm, for computing $\frac{\partial \hat{J}_1(W,b)}{\partial W_{ij}^{(l)}}$.

In this section, we derive an expression for $\frac{\partial \hat{J}_2(W,b)}{\partial W_{ij}^{(l)}}$ and furnish an algorithm for its computation. To proceed, let us recall from Section 3 that:

$$\hat{J}_2(W,b) \;=\; ||a^{(n_l)[t]}||_1 \tag{22}$$

Now, keeping the notation of Section 2 in mind, the above can also be written as:

$$\hat{J}_2(W,b) = \sum_{i=1}^{s_{n_l}} ||a_i^{(n_l)[t]}|| \tag{23}$$

In what follows, we will simply write $||a_i^{(n_l)}||_1$ instead of $||a_i^{(n_l)[t]}||_1$, for a less clumsy notation, and the reader is implored to keep this in mind. The ultimate plan is to derive a variant of the backpropagation algorithm for the computation of $\frac{\partial \hat{J}_2(W,b)}{\partial W_{ij}^{(l)}}$ for all l, i and j. To this end, using the notation of Section 2, we first need to spell out the pertinent feed-forward equations:

$$z_i^{(l+1)} = \sum_{j=1}^{s_l} a_i^{(l)} w_{ij}^{(l)} + b_i^{(l)} \tag{24}$$

$$a_i^{(l)} = f(z_i^{(l)}) \tag{25}$$

where $f(.)$ denotes the relevant transfer function, which typically ally is the sigmoid, tan-sigmoid or linear transfer function, depending on the layer in question. Also, we define a 'sign function 'according to:

$$sign(a_i) \;=\; \begin{cases} +1 & \text{if } a_i > 0 \\ -1 & \text{if } a_i < 0 \\ 0 & \text{if } a_i = 0 \end{cases} \tag{26}$$

Based on the above, one can show, using the chain rule, that $\frac{\partial \hat{J}_2(W,b)}{\partial W_{ij}^{(n_l-1)}}$, which is the partial derivative with respect to an arbitrary weight that projects to the output layer, $l = n_l$, from the penultimate layer, $l = n_l - 1$, is given by:

$$\frac{\partial \hat{J}_2(W,b)}{\partial W_{ij}^{(n_l-1)}} \;=\; sign(a_i^{(n_l)}) f'(z_i^{(n_l)}) a_j^{(n_l-1)} \tag{27}$$

where $f'(.)$ denotes the derivative of the transfer function $f(.)$. In anticipation of what is to come next, we define a 'signed derivative', denoted $\beta_i^{(n_l)}$, as follows:

$$\beta_i^{(n_l)} \;=\; sign(a_i^{(n_l)}) f'(z_i^{(n_l)}) \tag{28}$$

Let us notice the similarity between the form of $\beta_i^{(n_l)}$ and the so-called scaled error, $\delta_i^{(n_l)} = (a_i^{(n_l)} - y_i) f'(z_i^{(n_l)})$ (where y_i denotes the i-th component of the current target output training vector), which plays a key role in the standard backpropagation algorithm. In particular, let us make the *key observation* that $sign(a_i^{(n_l)})$ in $\beta_i^{(n_l)}$ is playing the role which $(a_i^{(n_l)} - y_i)$ plays in $\delta_i^{(n_l)}$. Indeed, we can easily express $\dfrac{\partial \hat{J}_2(W,b)}{\partial W_{ij}^{(n_l-1)}}$ in terms of $\beta_i^{(n_l)}$ by putting Equation 28 into Equation 27:

$$\frac{\partial \hat{J}_2(W,b)}{\partial W_{ij}^{(n_l-1)}} \;=\; \beta_i^{(n_l)} a_j^{(n_l-1)} \tag{29}$$

To proceed, let us now move one layer back through the ANN, and attempt to write an expression for $\dfrac{\partial \hat{J}_2(W,b)}{\partial W_{ij}^{(n_l-2)}}$. Again, using the chain rule, one can show that:

$$\frac{\partial \hat{J}_2(W,b)}{\partial W_{ij}^{(n_l-2)}} \;=\; \sum_{p=1}^{s_{n_l}} sign(a_p^{(n_l)}) f'(z_p^{(n_l)}) W_{pi}^{(n_l-1)} f'(z_i^{(n_l-1)}) a_j^{(n_l-2)} \tag{30}$$

But, from Equation 28, it is clear that for any index variable, p, we can write $\beta_p^{(n_l)} = sign(a_p^{(n_l)}) f'(z_p^{(n_l)})$. Putting this into Equation 30, we readily obtain:

$$\frac{\partial \hat{J}_2(W,b)}{\partial W_{ij}^{(n_l-2)}} \;=\; \sum_{p=1}^{s_{n_l}} \beta_p^{(n_l)} W_{pi}^{(n_l-1)} f'(z_i^{(n_l-1)}) a_j^{(n_l-2)} \tag{31}$$

At this juncture, we point out that if $\hat{J}_2(W,b)$ had been defined by $\hat{J}_2(W,b) = \dfrac{1}{2}\sum_{i=1}^{s_{n_l}} (a_i^{(n_l)[t]} - y_i)^2$ rather than $\hat{J}_2(W,b) = \sum_{i=1}^{s_{n_l}} ||a_i^{(n_l)[t]}||$, then we would have had $\dfrac{\partial \hat{J}_2(W,b)}{\partial W_{ij}^{(n_l-1)}} \;=\; \delta_i^{(n_l)} a_j^{(n_l-1)}$ and $\dfrac{\partial \hat{J}_2(W,b)}{\partial W_{ij}^{(n_l-2)}} \;=\; \sum_{p=1}^{s_{n_l}} \delta_p^{(n_l)} W_{pi}^{(n_l-1)} f'(z_i^{(n_l-1)}) a_j^{(n_l-2)}$. This observation indicates that, for the problem at hand, we can obtain a backpropagation algorithm for computing the partial derivatives of $\hat{J}_2(W,b)$ simply by letting $\beta_i^{(l)}$ play the role which $\delta_i^{(l)}$ normally plays in the standard backpropagation algorithm. As a specific implication of the above observation, one could define $\beta_i^{(n_l-1)} \;=\; \sum_{p=1}^{s_{n_l}} \beta_p^{(n_l)} W_{pi}^{(n_l-1)} f'(z_i^{(n_l-1)})$, in analogy with what is usually done in the standard backpropagation algorithm setting, and then write $\dfrac{\partial \hat{J}_2(W,b)}{\partial W_{ij}^{(n_l-2)}} = \beta_i^{(n_l-1)} a_j^{(n_l-2)}$. Now generalizing, for any $l \in \{1, 2, \ldots, n_l - 1\}$, we can define $\beta_i^{(l)}$ according to:

$$\beta_i^{(l)} \;=\; \sum_{p=1}^{s_{n_l}} \beta_p^{(l+1)} W_{pi}^{(l)} f'(z_i^{(l)}) \tag{32}$$

where the computation of $\beta_i^{(n_l)}$ (i.e. the base step) is already given in Equation 31. Hence, we are led to propose the algorithm below for computing the partial derivatives of $\hat{J}_2(W,b)$ with respect to ANN weights, (and biases as well):

Algorithm 2 A Back-Propagation Algorithm for Calculating $\dfrac{\partial \hat{J}_2(W,b)}{\partial W_{ij}^{(l)}}$ and $\dfrac{\partial \hat{J}_2(W,b)}{\partial b_i^{(l)}}$

Inputs: The set of current ANN weights, $W_{ij}^{(l)}$, and biases, $b_i^{(l)}$, $l = 1, 2, \ldots, n_l$, $i = 1, 2, \ldots, s_l$, and $j = 1, 2, \ldots, s_{l-1}$, along with a specific input vector, $x^{[t]} = (x_1^{[t]}, x_2^{[t]}, \ldots, x_{s_1}^{[t]})$.

1). Perform initialisation by filling ANN 'input activations' with the input vector, $x^{[t]}$.
That is, in the input layer, FOR each $i = 1, 2, \ldots, s_1$, set:
$a_i^{(1)} := x_i^{[t]}$
END FOR

2). Perform a Feedforward pass through the ANN, computing all activations through out the network.

That is, FOR each $l = 1, 2, \ldots, n_l - 1$, for each $i = 1, 2, \ldots, s_{l+1}$, for each $j = 1, 2, \ldots, s_l$, set:
1a). $z_i^{(l+1)} := W_{ij}^{(l)} a_i^{(l)} + b_i^{(l)}$
1b). $a_i^{(l+1)} := f(z_i^{(l+1)})$
END FOR

3). At the output layer, (i.e. layer L_{n_l}), of the ANN, FOR each i compute:
$\beta_i^{(n_l)} := sign(a_i^{(n_l)}) f'(z_i^{(n_l)})$
END FOR

4). FOR each $l = 1, 2, \ldots, n_l - 1$, for each $i = 1, 2, \ldots, s_{l+1}$, set:
4a). $\dfrac{\partial \hat{J}_2(W,b)}{\partial W_{ij}^{(l)}} = \beta_i^{(l+1)} a_j^{(l)}$
4b). $\dfrac{\partial \hat{J}_2(W,b)}{\partial b_i^{(l)}} = \beta_i^{(l+1)}$
4c). $\beta_i^{(l)} = \sum_{j=1}^{s_{l+1}} \beta_j^{(l+1)} W_{ji}^{(l)} f'(z_i^{(l)})$
END FOR

---

6 A Gradient Descent Algorithm for Minimizing SENNS's Objective Function

In this section, we describe a gradient descent algorithm for the minimisation of our objective function, which we gave in Equation 4. The algorithm is labeled as Algorithm 3 below. The algorithm relies on Algorithms 1 and 2 for computing $\dfrac{\partial \hat{J}_1(W,b)}{\partial W_{ij}^{(l)}}$ and $\dfrac{\partial \hat{J}_2(W,b)}{\partial W_{ij}^{(l)}}$ respectively. It then implicitly utilizes

Equations 8, 9 of Section 3, for the computation of $\frac{\partial J_1(W,b)}{\partial W_{ij}^{(l)}}$ and $\frac{\partial J_2(W,b)}{\partial W_{ij}^{(l)}}$ respectively, before using Equation 7 of that same section to compute $\frac{\partial J(W,b)}{\partial W_{ij}^{(l)}}$. Finally, using a learning rate, α, the algorithm proceeds to update ANN weights in a gradient-descent fashion according to:

$$W_{ij}^{(l)new} = W_{ij}^{(l)old} - \alpha \frac{\partial J(W,b)}{\partial W_{ij}^{(l)}} \tag{33}$$

where $W_{ij}^{(l)old}$ and $W_{ij}^{(l)new}$ denote the old and new weights respectively. Without any further ado, here is the algorithm we set out to derive:

Algorithm 3 Minimisation of SENNS's Objective Function via Gradient Descent: Weights Version

Inputs: Regularisation parameters, $\lambda_1, \lambda_2, \lambda_3, \lambda_4$, a learning rate, α, a training set, $x^{[t]}, t = 1, 2, \ldots, m$, parameter $M_{\mathbb{C}}$, which is the number of training set pairs belonging to the same class, parameter $M_{\mathbb{D}}$, which is the number of training set pairs belonging to different classes, and a set of initial randomized small weights, $W_{ij}^{(l)}$, $l = 1, 2, \ldots, n_l$, $i = 1, 2, \ldots, s_l$, and $j = 1, 2, \ldots, s_{l-1}$.

1). Perform the pair of initialisations:

1a). $\dfrac{\partial J_1(W,b)}{\partial W_{ij}^{(l)}} := 0$

1b). $\Delta J_2 := 0$

2). FOR each pair $(x^{[t]}, x^{[u]}), t = 1, 2, \ldots, m, u = 1, 2, \ldots, m$

2a). Using Equation 3, set $\mathbb{S}(x^{[t]}, x^{[u]}) := \dfrac{\lambda_1}{M_{\mathbb{C}}}$

2b). Using the current pair, $(x^{[t]}, x^{[u]})$, compute $\dfrac{\partial \hat{J}_1(W,b)}{\partial W_{ij}^{(l)}}$ via Algorithm 1

2c). Increment $\dfrac{\partial J_1(W,b)}{\partial W_{ij}^{(l)}} := \dfrac{\partial J_1(W,b)}{\partial W_{ij}^{(l)}} + \mathbb{S}(x^{[t]}, x^{[u]}) \dfrac{\partial \hat{J}_1(W,b)}{\partial W_{ij}^{(l)}}$

END FOR

3). FOR each $x^{[t]}, t = 1, 2, \ldots, m$

3a). Compute $\dfrac{\partial \hat{J}_2(W,b)}{\partial W_{ij}^{(l)}}$ via Algorithm 2

3b). Increment $\Delta J_2 := \Delta J_2 + \dfrac{\partial \hat{J}_2(W,b)}{\partial W_{ij}^{(l)}}$

END FOR

4). Set $\dfrac{\partial J_2(W,b)}{\partial W_{ij}^{(l)}} := \dfrac{\lambda_3}{m} \Delta J_2$

5). Set $\dfrac{\partial J(W,b)}{\partial W_{ij}^{(l)}} := \dfrac{\partial J_1(W,b)}{\partial W_{ij}^{(l)}} + \dfrac{\lambda_3}{m} \dfrac{\partial J_2(W,b)}{\partial W_{ij}^{(l)}} + \lambda_4 W_{ij}^{(l)}$

6). Finally, update the ANN weights via gradient descent using learning rate α:

$$W_{ij}^{(l)} = W_{ij}^{(l)} - \alpha \left(\frac{\partial J(W,b)}{\partial W_{ij}^{(l)}} \right)$$

7). Repeat STEPs 1 to 6 untill convergence, or a maximum number of iterations is reached

7 Conclusion

In this concept paper, we have proposed a technique called SENNS (Sparse Extraction Neural NetworkS) for the feature extraction problem, which is a problem at the heart of pattern recognition and machine learning. Philosophically, our proposed method draws on the idea of extracting features that maximise inter-class variances, while minimising intra-class variances. As a result, the method fits immediately into the framework of graph embeddings. However, unlike two of the most representative members of the class of methods within the graph embedding school of thought, our proposed SENNS enforces sparsity on the extracted features, and utilises powerful non-linear projections rather than linear, kernel or tensor transformations, to effect the feature extraction process. We formulated SENNS as the minimisation of a regularised sum of four terms, and derived an effective gradient descent algorithm for the resulting minimisation problem. Via rigorous mathematical analysis, we showed how our algorithm can be specified as a set of tasks involving the standard back-propagation procedure, up to a modification for L_1 norms. Finally, a natural next line of action is to test SENNS out on some standard machine learning datasets such as the ARABASE database of Arabic characters, the CMU PIE database of faces, as well as the MNIST database of digits.

References

[1] Trier, O.D., Jain, A.K. and Taxt, T. Feature Extraction Methods for Character Recognition: A Survey. *Pattern Recognition*, vol 29, no 4, pp 641-662, 1996.

[2] Belkin M. and Niyogi, P. Laplacian Eigenmaps and Spectral Techniques for Embedding and Clustering. *Advances in Neural Information Processing System*, vol. 14, pp. 585-591, 2001.

[3] Etemad K. and Chellapa, R. Discriminant Analysis for Recognition of Human Face Images. *Journal of the Optical Society of America A*, vol. 14, no. 8 pp. 1724-1733, 1997.

[4] Turk, M. and Pentland, A. Face Recognition Using Eigenfaces, *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, pp. 586-591, 1991.

[5] Hu, M.K. Visual Pattern Recognition by Moment Invariants. *IRE Transactions on Information Theory*, vol. 8, pp 179-187, 1962.

[6] Chao K. and Srinath M.D. Invariant Character Recognition with Zernike and Orthogonal Fourier Mellin Moments. *Pattern Recognition* 35, 143-154, 2002.

[7] Lowe, D. Distinctive Image Features from Scale-Invariant Keypoints. *International Journal of Computer Vision* vol. 60, no. 2, pp. 91110, 2004.

[8] Belongie, S., Malik J., and Puzicha J., Shape Matching and Object Recognition using Shape Contexts. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2002.

[9] Blum, H. A. Transformation for Extracting New Descriptors of Shape. In: Models for the Perception of Speech and Visual Form, W. Wathen-Dunn, eds., *MIT Press*, Cambridge, MA, pp. 362-380, 1967.

[10] Kimia, B.B, Tannenbaum, A.R. and Zucker, S.W. Shape, Shocks and Deformations I: The Components of 2d Shape and the Reaction-Diffusion Space, *International Journal of Computer Vision* 15 189-224, 1995.

[11] Macrini, D., Dickinson, S., Fleet, D., and Siddiqi, K. Bone Graphs: Medial Shape Parsing and Abstraction, *Computer Vision and Image Understanding* 115, 10441061, 2011.

[12] Yan, S., Xu, D., Zhang, B., Zhang, H.-J., Yang, Q. and Lin, S. Graph Embedding and Extensions: A General Framework for Dimensionality Reduction. *IEEE Transactionson Pattern Analysis and Machine Intelligence*, vol. 29, no. 1, pp. 4051, 2007.

[13] Sundaramoorthi, G. and Yang, Y. Matching Through Features and Features Through Matching. Technical Report, KAUST.

[14] Dejiver, P.A. and Kitler, J. *Pattern Recognition: A Satistical Approach*. Prentice Hall, London, 1982.

[15] Ben Amara, N., Mazhoud, O., Bouzrara, N. and Ellouze, N. ARABASE: A Relational Database for Arabic OCR Systems. *International Arab Journal of Information Technology*, 2005.

[16] Sim, T., Baker, S., and Bsat, M. The CMU Pose, Illumination, and Expression Database. *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 25, no. 12, pp. 1615-1618, Dec. 2003.

[17] LeCun, Y., Bottou, L., Bengio, Y. and Haffner, P. Gradient-based Learning Applied to Document Recognition. *Proceedings of the IEEE*, 86(11):22782324, 1998.

[18] Ng A., Jiquan N., Chuan F., Yifan M. and Caroline S., *UFDL Tutorial on Neural Networks*. Web. August 2014 <http://ufdl.stanford.edu/wiki/index.php/Neural_Networks>

[19] Yang, J., Yu, K., Gong, Y. and Huang, T. Linear Spatial Pyramid Matching using Sparse Coding for Image Classification. In *IEEE International Conference on Computer Vision and Pattern Recognition*, 2009.

[20] Gao S., Tsang I., Chia L., and Zhao P. Local Features are not Lonely - Laplacian Sparse Coding for Image Classification. In *IEEE International Conference on Computer Vision and Pattern Recognition*, 2010.